The work on which the content of this book is centred took place over more than a decade. It started, in the middle fifties, with the construction of adaptive training machines, with superficially disconnected studies of chemical computing systems, and, towards 1960, with experiments on machine-monitored small group inter-action. Since that period an underlying theory has emerged from a gaggle of prescient concepts. It owes a great deal of its present shape to the ideation and criticism of friends and colleagues, only some of whom can be mentioned directly.

First of all, it is noteworthy that parallel work has gone on in two places; my own laboratory at System Research Ltd and in Heinz Von Foerster's Biological Computer laboratory, at the University of Illinois. Both endeavours were encouraged by Warren McCulloch; the reader will detect the influence of his ideas and guidance appearing repeatedly throughout the discussion. Apart from this, the parallel development was not specially contrived and it was sustained by irregular personal liason. Hence, it is gratifying to find that recent publications from the Biological Computer Laboratory image our own conclusions, differing, chiefly, in the notation employed and the area of application. People familiar with the field will probably find the threads of mutualism quite obvious; for the benefit of others, a few of these threads are picked out. For example, Loefgren worked with Von Foerster whilst refining the formalism on which the currently-used type of abstract reproductive and evolutionary process is founded; Matturana (whose theory of autopoietic systems is the analogue, in a biologist's mind, for certain stable cognitive organisations in the present theory) worked there as well; Matturana's theory is to appear in a subsequent monograph in this series. Both Ashby, the system theorist, and Gunther, the philosopher, taught and researched with Von Foerster; much of the present theory hinges upon their ideas.

On home ground, the theory, and the experiments as well, owe a great deal to two colleagues of long standing: Brian Lewis and Bernard Scott. Prof. Lewis and I shared a common interest in cognition since early in the 1960s and brooded jointly (as we still do) over problems of learning and teaching. Bernard Scott came to the laboratory at the time when Lewis went off to study education in the large and since that time we have maintained a comparably symbiotic intellectual relationship. At about the time I started to write this manuscript (having discarded many previous drafts of it as inadequate) both circumstance and research interest brought all of us into close contact again and we were joined, in the last year, by Dionysious Kallikourdis (who contributes one Appendix, explicitly).

In addition, I am indebted to my graduate students at Brunel University (in particular, Messrs. Abel, Ben Eli, Glanville, Perriera, Robinson, and Stavrinides); to faculty members at the Open University (notably Prof. Lewis, Prof. Neil, and Mr MacDonald Ross); to the faculty at the School of Information and Computer Science in the Georgia Institute of Technology (to Prof. Seigmann, Prof. Vallach and Mr. Barralt Torrijos for especially long and painstaking discussion); to Xavier Salazar Resines and other faculty members in the Institute of Engineering, U.N.A.M., Mexico: to Felix Kopstein and later to Bob Seidal (of HUMRRO) for clarifying many points in the psychology of computer based teaching. Warren Brodey and Gregory Bateson dissuaded me from the uncritical espousal of "goals"; Ivan Illich, John Lilley and George Spencer-Brown assured me, in very different ways, that cognition is unbounded, as is proposed by the present theory.

I would like to thank the patrons who supported and support this work both financially and by way of guidance: the Social Science Research Council of Great Britain* and the A.F.O.S.R. through its European Office. In common with others in this field, we owe a special debt to Prof. Rowena Swanson who insisted throughout the formative years between 1961 and 1967, when she served in that organisation, upon the proper communication and integration of ongoing research.

*Since mid 1974 the application and development of this work in education is a Social Science Research Council Research Programme at System Research.

In conclusion, several people (some already mentioned) have assisted at a different level. Prof. Lewis, by organising the manuscript and its author; Mr. MacDonald Ross by preparing a glossary, Mr. Scott, abetted by Miss Barsby, for collecting and checking the references and for making sense from a mass of near illegible illustrations.

Gordon Pask

Richmond, Surrey

March, 1973

Introduction

The theories and methods described in this book involve an approach to cognition which is coloured by interests and dispositions both of a philosophical and a pragmatic kind. The resulting bias is strong enough to be obtrusive and it is prudent to exhibit the idiosyncracies it engenders at the outset.

The theory is concerned with psychological, linguistic, epistemological, ethological, social or non-commitally mental events of which there is awareness. In general, awareness is manifest in the differentiated and relatively tractable form of consciousness: of something with someone. Awareness itself is a non-specific phenomenon, but consciousness is not: "A is conscious with B of R". For a concrete example, "I am conscious with you of the glasses, the Perrier water and the wine bottles which form a configuration on this table". For an abstract example, "I am conscious with you of a physical principle" (say, "the conservation of energy" or "Ohm's Law"). If awareness is not in evidence, but there are grounds for believing it might be, then the theory is concerned with why this should be so and is tailored to explain the curious fact that some mental processes are executed subconsciously or unconsciously; for instance, the conduct of a well learned skill, or most searches in long term memory.

The position adopted is compatible with phenomenology in general. But the theory's scope is narrowed insofar as certain forms of behaviour, commonly studied under the heading "psychology", are excluded from consideration. If a behaviour is accompanied by awareness then all is well; if not, especially if the method of observation is contrived to obfuscate or conceal all vestiges of awareness, then the behaviour is rechristened with (slightly obsessive) rectitude a "physiological process" and is placed beyond the theory's domain of interpretation; though not, of course, discounted because of that.

In contrast, the theory's scope is extended by an underlying contention that cognitive processes (involving actual or potential awareness) are not confined to the human brain. Though brains, human or animal, are often associated with cognitive operations, neither biological fabric nor any other kind of fabric is responsible (except in an incidental way) for the peculiar nature of cognition. Without contesting the utility and cogency of arguments from evolution (the gradual development of language, problem solving capability, awareness and the like) they tell, as they stand, a one sided tale about a ubiquitous and many faceted event. The present theory does not "biologicise" mankind or mind; nor is it reductionist. Cognition may occur at the level of groups of people (so social awareness is taken in earnest, not as a collation of individual awarenesses) or it may characterise the activity of slightly unconventional computing machines.

In this connection, there is a tacit and non-standard philosophical commitment regarding the inferences to be made if machine cognition *is* demonstrated; because of which the theory does not naively "mechanise" mankind or mind, either. The commitment is most clearly exhibited in the context of studies in "artificial intelligence".

As a rule, such essays are taken as attempts to imitate mentation, and as a popular corollary, to degrade uniquely mind-like processes, into no more than mechanical operations. The present theory is out of kilter with the mooted objective as well as the mooted value judgment. If a machine-executed intellect exists, according to a criterion related to but not identical with Turing's test (the details of which are later delineated), then the system in question is more general than a biological system. In particular, "artificial" is a misplaced qualifier. The proper term is "general intelligence".

This generality gives rise to some important, though often over-looked, practical consequences. Recall, first, that cognition has a wider connotation than usual; which is, however, captured by the posited test for intellect. Under this proviso, if cognitive processes can be realised in a general machine then it is possible to execute mental operations in artifacts that are not necessarily subject to the embarrassing spatio-temporal limitations and structural frailties of a biological processor. By this means it may be possible to secure a modicum of intellectual continuity, essential to social and technical development, which is nowadays approximated by ad hoc and suboptimal expedients. This possibility surely justifies efforts to realise general intelligence, whether or not more local objectives (pattern recognition, robotry for special purposes, and so on) are deemed valuable in their own right.

As orienting comments, I conjecture that generalised intellect is possible and that it does not preclude those affective components of mind that are often brushed aside whe machinery is brought into the picture. I definitely maintain that a substantial part of the human mind can be encoded for execution in machines more general than a human brain. By way of summary, cognition in the present and broad sense of the word, is a general process. Biological cognition is a rich and important, but nevertheless, very specialised, example of it.

Whether or not a theory with the orientation just outlined will prove attractive to a potential user depends entirely upon the sort of question the user wants to pose and the kind of explanation he regards as a satisfactory answer to it. For example, if a user is anxious to predict gross rates of adaptation, then the theory is unduly cumbersome and several statistical learning theories are more suitable vehicles for the enquiry. By the same token, if the user is testing a specific neurophysiological or psychophysical hypothesis, for example, about perception or about filtering operations, then the present theory is either inappropriate or frankly inapplicable. In contrast, the present theory comes into its own if the user's primary interest is focussed upon styles, methods, and strategies of learning, communication and thinking, upon the performance of complex skills, upon the structure of knowables, and so on. The theory is also attuned to answering such global questions as "why do people direct their attention (usually, to one goal at once)?" or "why are there distinct sensory and descriptive modalities?".

Under these circumstances, the theory is a serious candidate for adoption as the preferred explanatory or predictive device, provided the user is looking for the kind of systemic or Cybernetic explanation it is designed to furnish, rather than an explanation tied to particular assumptions about brains or configurations of input and output paths. The last two question topics will be used to illustrate the sort of explanation this is; at least to the extent of exhibiting some salient and distinctive features.

The first illustrative question, "why do people direct their attention.......?" is usually answered in terms of processor limitations. The cognitive process is housed in a fixed processor (the brain) with restricted storage capacity and execution rate. The argument proceeds, with some variations, as follows. At any instant, the processor, constrained by the mind like programme under execution is in some state. To satisfy the restrictions (of storage capacity etc.), the processor states, induced by executing the process, are partitioned, thus imposing an hierarchical structure on the processor. At any instant, higher level components switch in one of a limited number of lower level subprocessors. These subprocessors have the status of physiological components of any brain: they are the filters (Treisman) or the sensory analysers (Sutherland) of perception. If the sensory channel is disregarded (so that chunks of data, encoded from any modality, are under discussion) these conditions imply that all the data chunks refer to one goal (meaning, they are under one locus of control). In this case, the hierarchical structure is manifest as a series of functional rather than physiological demarcations; for example, the demarcation of short term, intermediate, and long term storage. The rigid model is relaxed, to match reality, by introducing a small "leakage" of information around the hierarchical network (to allow for the shift of attention), by introducing special buffers in the visual modality, and so on.

None of this is denied. In particular, there is no question of raising pedantic objections to the waivers that match the rigid model (a parody in any case) to psychological reality. Factors of the type mooted in the last paragraph may, very often, impose the limits that do canalise and restrict the field of attention. But I find this type of explanation unsatisfactory, even so, because it depends upon constraints of limited tenure which taken in wider compass, prove arbitrary. Hence, the present theory is designed to explain attention span and attention directing as a necessary property of the cognitive process in a processor-independent manner; that is, in a way that is applicable to any processor in which the cognitive process could be executed. The explanation goes along these lines.

Any cognitive *process*, later identified with a *concept*, is shown to reproduce or, equisignificantly, to stabilise a relation. Generally this is a relation between entities in an environment, though it need not be. Further, it is shown that any non-transient process (an observable concept) is reproducible in the context of a certain class of relations. The word "reproducible" is used in the sense of the theory of abstract (or symbolic) reproductive automata, onto which most cognitive activities are mapped. Under execution, a concept (alias process) owes its

integrity to actual reproduction, not just to the fact that it is reproducible. As a result, an appropriate context must exist and, by hypothesis, this together with the process under execution, is the field of attention. The argument proceeds by noting that any field of attention is restricted for the following reasons. (a) Execution and reproduction of concepts (in a given context) gives rise to an "immune response" by which a process recognises copies of itself and repudiates sufficiently dissimilar processes. (b) In aggregate, reproduction and exclusion leads to a homeostatic trapping condition (later called "cognitive fixity") that narrows and momentarily stabilises the field of attention. (c) As a consequence of the same dynamic principles, the trapping condition is transient; the field of attention must either change or be changed, perhaps by a change in context, perhaps by a change in the concept under execution. (d) The fresh field of attention will, in turn, be narrowed and stabilised.

The reader may rest assured that this does not count as an "explanation"; it is, at most, a sketch for the type of explanation to be offered. On the other hand it should be noted that a properly constructed explanation of this type is processor-independent; for example it applies just as well to activity in one of Aleksander's (1973) dynamic meshes as it does to the cognitive operations executed in a brain. Moreover, the existence of intermediary constructs such as "dissimilar processes" and "cognitive fixity" can be empirically verified. They are manifest as distinctive and (in the same context) exclusive learning strategies, as various inertial phenomena, and so on.

The second illustrative question "why are there different sensory or descriptive modalities.......?" is frequently answered by the terse reply "Because people have eyes and ears".

There is no denying it. They do. However the answer is unsatisfying because it fails to indicate whether or not any cognitive process must be associated with equivalent mechanisms; either by way of sensors like eyes and ears that are attached to the processor in which the process is executed or as modes of representation, like "visual imagery" or "verbal imagery" that constitute part of the programme organisation.

The present theory answers the second illustrative question by affirming the necessary existence of modes whenever cognition is observable. Under circumstances that are to be delineated, the replication of a cognitive process (for example, the replication of a concept of R in the earlier paradigm "A is conscious with B of R") depends upon the existence of a conversation between distinct and integral participants A and B. The reproduction of the process is manifest in the A, B, dialogue as a complex of explanations, from A to B and vice versa, which is later called an *understanding*. The crucial point is that replicative cycles can be observed and manipulated if and only if they are reified in terms of understandings that split the cycle into parts which belong to A and to B. (It should be emphasised that A's understanding of B and B's understanding of A, which is the event detected in the already mentioned extension of Turing's test, has no direct connection with an observer's understanding of the A, B, dialogue).

What are A, and B, in this formulation? One possibility is that A and B are individuals engaged in discourse; if so, they each have potentially private descriptions of R (namely A's description and B's description) a common part of which is shared if R is understood. Another possibility is that A and B represent organisational entities such as the "proposer" and the "critic" routines in an intelligent program (the importance of which is stressed, by Marvin Minsky). The present theory admits an indefinite number of interpretations of this kind including, for example, the following: (a) The entities A and B are parts of a mind, executed in a common brain, that have different modes of representation for the same relation (visual and verbal, for instance) or (b) A is the organisation of a visual apparatus attached to the brain qua processor; B is the organisation of an auditory apparatus attached to the same processor. In any case, an observation of understanding (hence, also, an observation of a viable cognitive system) relies upon an A, B, distinction. Sensory modalities or modes of representation that happen to exist are special cases of this distinction which must, somehow, be realised if mentation is in evidence.

These remarks convey the flavour of the theoretical development and give some idea of its scope. Very similar theories have been advanced in the "hard" sciences, notably physics and genetics. Regarding the development as primarily psychological, it leads to a reappraisal of such notions as concept, memory, individual, learning, creativity and innovation. Such constructs are (at least) restated in terms which accord more closely with ordinary-language usage than the often stilted connotations of

the traditional discipline (for example, the obviously unsatisfactory, though locally useful, meaning of "concept" as a kind of "class" in storage).

Viewed from the wider perspective of cognition unqualified, the theoretical development is consonant with and dependent upon fairly recent advances in computation and information science, which belong to five major categories: (1) The theory of reproductive and evolving automata. This theory, which has blossomed after a lengthy gestation, also underpins the thesis advanced by Matturana. (2) Formalisations for parallel and (distinctly) concurrent processes, together with the closely related notions of a "non-deterministic program" and a "fuzzy algorithm", otherwise called "heuristics"[1]. Though serial execution is possible (entailing, however, special resolution mechanisms operating at each step) such programs/algorithms are generally/unrestrictively executed as parallel processes. Further, their execution may give rise to concurrency, insofar as otherwise asynchronous loci of control are brought into local synchronicity by virtue of interaction. (3) The series of distinctions, mainly due to Petri and Holt, between control and information transfer. The latter event takes place whenever previously asynchronous processes become synchronised (this usage, which is highlighted when it is employed, is clearly distinct from "selective information" estimated by someone who is observing the system). (4) Widely dispersed work on search and inference, mostly in relational networks and data structures richer than trees. (5) Recent results on program or process equivalence (used informally in this volume, more seriously in the next volume).

The theory also begets several philosophical points of contention, of which the following have special interest. (a) The popular interpretation of a "goal" as an "end state" is replaced by an interpretation of "goal" as an intention. This, amongst other things, permits the consideration of underspecified goals which are, in fact, more common. (b) The languages used to mediate the conversations (for example, between A and B about

1 The identification between "Fuzzy Algorithms" and "heuristics" is also due to Roger Hartley (I.D.E.A., Brunel University, 1973). I am not certain which, if either, of us has precedence. Nor does it matter a great deal provided that the much used but much maligned word "heuristic" is given an acceptable meaning.

R) which figure prominently in the theory are command and question languages in which personal pronominalisation (addressing) is permissible, as well as the impersonal pronominalisation of objects. Because of that, the theory has a reflective component; and whether it is reflective or not, the entire theory is relativistic. (c) Predication or distinction is adumbrated by the theory (rather than placed outside the theory, as it is by giving a universe of objects and giving the properties in terms of which they are described). (d) In contrast to languages in which the basic statements are propositions designating simple relations between objects or classes, the basic statements in the conversational languages are metaphors, designating material analogies that always refer to the participants.

The theory has been presented, criticised, and revised on several occasions; both by colleagues and myself. As a result, many potential objections have been eliminated. The major remaining criticism, probably due to the emphasis upon processes rather than processors, is that we are fiddling with disembodied minds. This objection stems from crossed habits of thinking; not from a matter of fact. Hence, it cannot be legitimately circumvented and it would be dishonest to gloss the occasionally unwelcome revisions of attitude that are entailed. The truth is, we might be playing tricks with pure mentation, as the critics fear. Quite conceivably, that would be an interesting exercise in its own right (and it is no more nor less justifiable than the traditional pursuit of disenminded bodies). In fact, we are chiefly concerned with cognition that goes on; as a result, with cognition that goes on in some processor. But it is still true that the processors in question are generally non specific; that one process may be executed in several processors and that several processes may be executed in one processor.

The remaining idiosyncracies of the book are due to the chosen field of application rather than anything inherent in the theory itself.

Most of the work has been done and most of the data has been obtained in connection with complex skill learning and education. Useful investigations, in the latter area especially, depend upon lengthy and individualised experiments in which students come to grips with sizeable bodies of subject matter and build non trivial cognitive representations of whoever (or whatever) occupies the role of teacher. Hence, the experiments

to be described bear little resemblance to studies in the psychological laboratory which commonly employ miniature learning situations. With few exceptions, we deal with learning which is, by conventional standards, unmanageably complex. As a result, the "experiments" either are or are "near to" real life applications (in teaching, course design, and so on). It can be argued that, because of this, the data is likely to be uninterpretable. But that argument, though superficially reasonable, is categorically denied. On the contrary, we maintain that the canon of simplicity is frequently misapplied: the situation which seems intuitively simple often turns out to be actually complex, if only because a simple task occupies only part of the mind (the other, essentially uncontrolled, mental activities being responsible for most of the variance in the data). It is certainly true that experiments with situations that approximate real life learning call for the use of special methods and techniques. But these methods stem in a very direct way from the underlying theory and, if they are employed, the resulting data is clear cut and usually unequivocal. If this contention is adequately buttressed (as I believe it is, by the results to be described), then it lends strength, or at least pragmatic credibility, to the theory from which the methods are derived.

It is generally agreed that "educationally relevant" or "realistically large" learning and teaching processes ought to be examined. It is also conceded that most investigations of learning singularly fail to produce the type of information which is useful to an educator: the laboratory studies, because the conclusions reached cannot (for the reasons suggested in the last paragraph) be summed to produce a picture of mentation as a whole; the field studies, because there is no structure in which to frame hypotheses, to make measurements, and to integrate results.

In the absence of such a structure, the investigation of educationally realistic learning degenerates into data collection; either haphazard or guided by a system having administrative rather than scientific validity. Due to the sheer volume of data that emanates from the enquiry, the results must somehow be summarised before they are inspected. Various gambits are possible; for example, aggregation by head counting or the application of statistics designed to compare gross (and often group averaged) indices of performance before and after a learning experience. Without serious exception, these operations obscure any information about individual differences, the style and the form of learning, that might have been garnered from the original records. We maintain (and neither practising teachers nor thoughtful educational theorists seem to disagree) that this lost or unobtainable information is crucial in recommending means for catalysing and regulating the learning process, for optimising the organisation of subject matter or achieving similar desirable ends.

Hitherto, it was only possible to reconstruct a small part of the missing information by augmenting the gross (field study) results with principles of learning derived from laboratory experiments. The composition was both improvident and imperfect. The present theory and the methods derived from it are claimed to obviate this expedient by providing a scientifically valid structure for studies of large scale learning at an individual level and thus to resolve the dilemmas begotten by the lack of such a thing. It is believed (with some empirical support) that if the theory and methodology are applied to other intractable areas of enquiry (for example, opinion sampling, social dynamics, or market research) then a similar advantage is obtainable. These applications are chiefly considered in the next volume.

The book is organised as follows:

Chapter 1 introduces a basic observational paradigm, the *steady state* technique, in the context of perceptual motor learning, where it is most readily exhibited. Chapter 2 develops the idea of *relativistic* observation and lays the foundations for a molar or macro-grain theory of learning and cognition, based upon the minimal unit of a conversation. Chapter 3 is concerned with a further experimental method, the *cooperative externalisation technique* (or CET) employed to exhibit usually private cognitive processes as part of an observable dialogue. Chapter 4 is devoted to a description of a partially or wholly mechanised experimental facility, CASTE: (Course Assembly System and Tutorial Environment) which is used to foster transactions of a type to be characterised as *strict conversations*, that are anchored on a specific subject matter and punctuated by occasions upon which the participants achieve *understanding*. Apart from its role as an experimental facility, CASTE is also a *model* for the molecular or micro-grain theory of learning and

cognition, concerned with concepts, cognitive interactions and so on. At this point, however, only two theoretical constructions are seriously considered; namely a heuristic for executing the CET which establishes the least biassed kind of strict conversation; and a further heuristic, the uncertainty regulation heuristic, which modulates the unbridled transactions of the CET so that they satisfy boundary conditions stemming from the molar or macro-grain theory and are held, at this level, to encourage rapid and effective learning. Chapter 5 contains the first proper formulation for a theory of strict conversations (notably, tutorial conversations in CASTE or their analogues in general discourse) and of the participating individuals (recall that the theory comments, quite radically, upon the status of participants). Chapter 6 retrieves the two special constructions of Chapter 4 as a precise statement; it is shown that a CET heuristic is a type of mirror: a *cognitive reflector* in which the student sees himself as though he were his own teacher. The uncertainty regulation heuristic is a mirror distorted by constraints that are specified by an external observer. Further, in this Chapter we begin the task of identifying all pertinent paradigms of psychology with suitable approximations to a strict conversation. Chapter 7 deals with the (by this stage) obtrusive requirement of building a conversational domain: the set of knowable relations on which the conversation is anchored. It is assumed (in this Chapter, only) that whatever may be known can be stated as a thesis by a subject matter expert who is free to say what he likes, provided the topics he moots are compatible with certain restrictions upon knowability or learnability. In Chapter 8 we justify this informal constructive operation by showing that it is identical with the evolutionary or course assembly operation-mode of CASTE. Here, the conversational domain evolves either under the guidance of a subject matter expert who replaces the student or under the guidance of a student who opts, from time to time, to adopt the role of innovator. Chapter 9 completes this part of the story by describing the other aspect of a conversational domain; what may be done (in contrast to what may be known); and by giving an account of the laboratory-like modelling facilities in which modelling operations, carried out by a student, are interpreted as non-verbal explanations. Though task specifications and modelling facilities are discussed quite generally, each salient example is referred back to the facility used in (the illustrative application of) the CASTE equipment. Chapter 10 contains a discussion of broader, but actually realised, applications of the theory; mostly in education. Chapter 11 ties up the threads of the argument, some of them in a speculative fashion. We deal with the intimate relation between the (previously separate) macro-grain and micro-grain theories and, as a result, reach some conclusions about self-reference (the reflective theory) and the conditions for consciousness. In addition, we review some of the topics to be discussed in the next volume.

Materials that are not essential to the main line of the argument have been relegated to Appendices, in order to allow for the concise statement of fairly complex and closely knit ideas. The reader is urged, however, to regard the Appendices as part of the story; they are not just depositories for optional reading matter. For preference, they should be scanned when they are referenced, or shortly afterwards. If that is done, nothing will be lost by perusing them at leisure.

1. Observation and regulation

Any kind of measurement, regulation, or control calls for a standard experimental condition in which observations are made and parameters of the system are adjusted.

The classical paradigm was spawned in the physical sciences (though in that province it now has only the limited status of a paradigm for observing coarse grained phenomena). The standard condition is a fixed frame of reference, an observer of which is impartial and external to the experiment. In most cases he is omniscient; all possible outcomes (those that are countenanced as relevant to the enquiry) are known beforehand and the observer entertains a definite hypothesis, either deterministic or probabilistic in calibre. The system to be observed is demarcated unequivocably as an *it* under scrutiny and is often divided into parts such as transmitter and receiver or causative input and output effect.

Commonly, one of these parts, the *environment*, is designed by the experimenter/external observer to react in a particular way to the object/assembly under examination. Failing that, the experimenter acts himself as the environment but his actions, in this capacity, obey strict and predetermined rules. Apart from these interactions, usually represented as values of independent variables, all other fluctuations (deemed irrelevant) are minimised, i.e. the experimental conditions are held constant, at any rate to the extent that they might influence the values of observed or dependent variables. If the experiment is repeated, every effort is made to replicate the same conditions. Finally, for several reasons (one of them the reductionist ethos underlying this method, another the ease of replicating particular entities) there is a tendency to carry out experiments that involve discrete perturbations (or the variation of one independent variable at once) and the observation of simple reactions (or few dependent variables).

1.1. *Behaviouristic formulation.* The classical paradigm is carried over, more or less piece-meal, into behaviouristic psychology though the nomenclature is changed. For example, to secure